

Abstract

A scheduling apparatus flexibly integrates guaranteed-bandwidth (GB) and best-effort (BE) flows and comprises a combination of a primary weighted-round-robin (WRR) scheduler (PWS) and a secondary WRR scheduler (SWS). The PWS distributes service to the individual GB flows and determines the amount of service that the BE flow aggregate should receive during each frame. The SWS takes care of fairly distributing the service share of the BE aggregate over the individual BE flows. The scheduling apparatus divides the service frame in two subframes. In the first subframe, the PWS fulfills the bandwidth requirements of the GB flows. In the second subframe, the SWS distributes fair service to the BE flows. For each frame, the duration depends on the amount of bandwidth allocated to the GB flows and on the number of GB flows that are backlogged at the beginning of the frame. The amount of bandwidth globally available to BE flows (i.e., the duration of the second subframe) is dynamically adapted to the backlog state of the GB flows, increasing when the GB-flow activity is low, and decreasing when the GB-flow activity intensifies.